

PATENT

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
(Attorney Docket No. 99CV02)

TITLE

**MODULAR BACKUP AND RETRIEVAL SYSTEM
USED IN CONJUNCTION WITH A STORAGE AREA NETWORK**

INVENTORS

John Crescenti
1 Ivy Road
Freehold, NJ 07728

David Alan Oshinsky
22 Francis Road
East Brunswick, NJ 08816

Srinivas Kavuri
40 Maple Court
Highland Park, NJ 08904

Anand Prahlad
3 Bucknell Drive
East Brunswick, NJ 08816


ASSIGNEE

CommVault Systems, Inc.
2 Crescent Place
Oceanport, N.J. 07757-0090

CERTIFICATE OF EXPRESS MAILING

I hereby certify that this correspondence, which includes 24 pages of Specification and 3 pages of Drawings, is being deposited with the United States Postal Service "Express Mail Post Office to addressee" Service under 37 C.F.R. Sec. 1.10 addressed to: Box Provisional Application, Assistant Commissioner for Patents, Washington, D.C. 20231, on July 6, 2000.

Express Mailing Label No.: EK 697087922 US


Teresa A. Newby

TITLE

**MODULAR BACKUP AND RETRIEVAL SYSTEM
USED IN CONJUNCTION WITH A STORAGE AREA NETWORK**

SPECIFICATION

~~CROSS-REFERENCE TO RELATED APPLICATION~~

~~This application claims the benefit of U.S. Provisional Patent Application Serial No. 60/143,743, filed July 14, 1999, pending.~~

~~INCORPORATION BY REFERENCE~~

This application hereby incorporates by reference, in its entirety, U.S. Provisional Patent Application Serial No. 60/179,345, filed January 31, 2000, and U.S. Provisional Patent Application Serial No. 60/143,744, filed July 14, 1999, both pending.

BACKGROUND

1. Technical Field.

The present invention is directed towards backup systems for computer networks. In particular, the present invention is directed towards the implementation of a distributed, hierarchical backup system with a storage area network (SAN) system.

2. Related Art.

Conventional backup devices commonly employ a monolithic backup and retrieval system servicing a single server with attached storage devices. These systems usually control all aspects of a data backup or retrieval, including timing the backup, directing the files to be backed up, directing the mode of the archival request, and directing the storage process itself through attached library media. Further, these backup and retrieval systems are not scalable, and often

direct only one type of backup and retrieval system, such as a network backup or a single machine backup.

Due to the monolithic structure of these backup and retrieval systems, a slight change in the operation of any one of the several functional aspects of the backup and retrieval system requires a large amount of effort to upgrade or change the backup and retrieval system, including in some situations, reinstalling the backup and retrieval system in its entirety.

Also, the operation of a backup and retrieval system across a network containing several different types of hardware and operating systems presents significant challenges to an enterprise scale backup including maintaining data coherency, bridging file system protocols, and accessibility issues across multiple hardware and operating system configurations.

Other currently available backup solutions do not address scalability issues, hierarchy issues, and the problems inherent in the storage of different files in a network file system. Many other problems and disadvantages of the prior art will become apparent to one skilled in the art after comparing such prior art with the present invention as described herein.

SUMMARY OF THE INVENTION

A file processor manages data transmission in a computer storage system. The file processor operates as a part of a computing system and may be implemented as programs running on a computational device. A management component module and at least one client component work in conjunction with the file processor for archival purposes such as archival requests. The client component may be implemented as a program running on a computing device. Archival requests include storing data such as a computer file in a location different than the original location of the data. Archival requests may also include retrieval of stored data and may include restoring data to a previous state such as retrieving earlier versions of a file. The computer storage system may be comprised of a media component and a client component that manage functions associated with a backup of a computer storage system.

Another aspect of the invention includes a modular network storage system in which a file processor directs the functions associated with the archival of data over a network. A plurality of backup devices, each having space for the archival of data are directed by a plurality of media components. Each media component is a part of a computing device and is communicatively coupled to one or more of the plurality of the backup devices and the file processor for controlling archival functions of the backup devices in accordance with the direction from the file processor. A plurality of client components each generate archival type requests to the file processor which then provide direction to the plurality of media components for directing the archival functions in accordance with the archival type requests.

The modular network storage system may include a management component that is communicatively coupled to the file processor and the plurality of client devices for coordinating

archival functions where the management component is a part of a computing device such as a program running on a computer. The modular network storage system may include a plurality of client devices where each client component is communicatively coupled to one or more of the plurality of client devices and the file processor for communicating the archival type requests from the client devices to the file processor. At least two of the plurality of client devices may run different operating systems. A network storage media may be communicatively coupled to two or more of the plurality of client devices over the network as well as the plurality of backup devices and at least one client device may include a local storage media, wherein the archival functions include reading data from either the network storage media or the local storage media and then writing the data to one of the plurality of backup devices.

A method of the present invention includes providing a file processor, which is communicatively coupled to at least one client component and a plurality of media components; providing a plurality of backup devices, each backup device has physical storage space for performing archival functions; coupling the plurality of media components communicatively with the plurality of backup devices, and with a file processor, wherein each of the media components control the archival functions of one or more backup devices; generating an archival type request, by the client component to the file processor; and directing, by the file processor through the plurality of media components, the backup devices to perform an archival function, in accordance with the archival type request.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a schematic block diagram of a modular backup and retrieval system built in accordance with principles according to the present invention.

Fig. 2 is a schematic block diagram of a modular backup system working in conjunction with a storage area network (SAN) system according to principles of the present invention.

Fig. 3 is schematic block diagram of the interaction of the library media of Fig. 2 with the SAN system.

044463.0024 AUSTIN 188229 v2

DETAILED DESCRIPTION OF THE DRAWINGS

Fig. 1 is a schematic block diagram of a modular backup system. A modular backup system 100 comprises three components, a management component 110, one or more client components 120, and one or more media components 130.

Typically, these three components, the management component 110, the client component 120, and the media component 130, may reside on several different machines. For example, the management component 110, the client component 120, and the media component 130 may all reside on a single computing device. Or, the management component 110 and one of the media components 130 may reside on a single computing device with the client component 120 residing on a different computing device. Or, the management component 110 and one of the client components 120 may reside on a single computing device with the media component 130 residing on a different computing device. Or, the media component 130 and the client component 120 may reside on the same computing device with the management component 110 residing on a different computing device. Or, the management component 110, the client component 120, and the media component 130 may all reside on different computing devices.

As shown in Fig. 1, the management component 110 is coupled to the client components 120 and the media components 130. The media components 130 are also coupled to the client components 120.

These components of the management component 110, the client component 120, and the media component 130 are typically software programs running on the respective computing

devices. Although the computing devices may not be the same devices, communication should exist between these components, as is demonstrated.

The client component 120 controls the actions and parameters of a backup or retrieval for a particular client computing device. A client computing device is the computing device in need of backup and retrieval assistance. The client components 120 each reside on a client computing device, or are in active communication with the client computing device. The particular client component 120 provides, for a particular client computing device, communication with a management director component 110 regarding such parameters as backup schedules, types of files in the backup schedule, the method of backup or retrieval, and other broad scope backup and retrieval management functions for the client computing device. The particular client component 120 communicates with a particular media component 130 responsible for the actual backup or retrieval function.

The media component 130 controls the actions and parameters of the actual physical level backup or retrieval at the library media containing the archived data. Each media component 130 is responsible for one or more physical backup media devices. As shown in Fig. 1, the media component 130 may be responsible for a single backup device 140, or for a plurality of backup devices 150 through 160. The particular media component 130 directs the data that is the subject of an archival type request to or from, as the case may be, the particular backup devices 140, 150, or 160 that it is responsible for. In the case of a retrieval type archival request, the particular media component 130 directs the retrieved data to a requesting client component 120.

The particular media component 130 also creates a library index for the data contained on the particular backup devices 140, 150, or 160 for which it is responsible for operating. Additionally, the particular media component 130 indexes the location of the archived data and files on the particular associated backup media devices 140, 150, or 160 that it is responsible for operating, and allows the management component 110 and the client component 120 access to certain information about the index entries. The media component 130 uses this library index to quickly and easily locate a particular backed up file or other piece of data on the physical devices at its disposal.

The particular media component 130 either resides on a computing device physically responsible for the operating the library media which the particular media component is responsible for, or it must be in active communication with that computing device. The media component also communicates with the management component 110, since the management component is responsible for the allocation of physical media for backup and retrieval purposes.

The backup devices 140, 150, and 160 can comprise many different types of media, such as massively parallel fast access magnetic media, tape jukebox media, or optical jukebox media devices. The determination of which backup device is to be implemented is determined by several parameters. These include time related frequency of accesses, importance of the backup file or data and urgency of its retrieval, or how long ago the backup was made.

The management component 110 directs many aspects of the backup and retrieval functions. These aspects include scheduling policies, aging policies, index pruning policies, drive cleaning policies, configuration information, keeping track of all running and waiting jobs, allocation of drives, type of backup (i.e. full, incremental, or differential), tracking different

applications running on each client, and tracking media. The management component 110 may contain the scheduling information for a timetable of backups for the computing devices. Any number of computing devices might be involved, and the computing devices may be interconnected.

Fig. 2 is a schematic block diagram of a modular backup system working in conjunction with a storage area network (SAN) system 250. A computing device 200 contains and operates a management component 202, which is responsible for the coordination of backup, storage, retrieval, and restoration of files and data on a computer network system 290. The management component 202 coordinates the aspects of these functions with a client component 212, running on another computing device 210, and a client component 222 running on yet another computing device 220. The computing device 220 also has an attached data storage device 214, to which it can store data and files locally.

The computing devices 210, 220, and 230 are connected to the SAN system 250 via a connection 264, such as a direct fiber channel connection, or a SCSI connection. However, it should be realized that any type of network connection is possible.

The SAN system 250 environment comprises the connection media 264, routers, and associated hubs for the actual data communication functions of the network, and a file processor 252. The elements of the SAN system 250 not explicitly numbered are implied in a remainder of the SAN system 250.

Another computing device 230 contains another client component 232. However, the computing device 230 is connected, through a network 270, to a file processor 252 for

interaction with the SAN system 250 through another network 265. This network could be any type of network, such as a LAN operating under a TCP/IP protocol.

The client components 232, 222, or 212 coordinate and direct local backup and retrieval functions on the computing devices 230, 220, and 210, respectively. The management component 202 coordinates and directs the overall network backup of the computer network 290.

The computing devices 210, 220, and 230 can all be different architectures of machines running different operating systems. Hardware systems could include those made by SUN, Hewlett/Packard, Intel based families of processors, and machines based on the RS6000 and PowerPC families of processors, to name a few. Operating systems can include the many flavors of UNIX and UNIX-like operating systems, such as HP/UX, Solaris, AIX, and Linux, to name a few, as well as Windows NT by Microsoft.

The file processor 252 of the SAN system 250 contains a client component 262 and a media component 260. Storage media 257, 258, and 259 are communicatively coupled to the file processor 252 for storage of network files from the computing devices 210, 220, and 230. These storage devices can be magnetic media for fast retrieval, tape media for longer term storage, or optical media for much longer term storage.

The overall SAN system 250 acts as a block access device to the computing devices 210, 220, and 230. Thus, the overall SAN system 250 acts as a virtual media device and centralizes the network file system from the computing devices 210, 220, and 230. As such, true dynamic sharing of the data and files through the SAN system 250 is possible. These data and files are available to the computing devices 210, 220, and 230. The computing devices 210, 220, and 230 present their network file and data requests to the file processor 252 over the SAN network

media 264 remainder of the SAN system 250 as they would any other storage media available to that computing device. The file processor 252, working in accordance with its software, interprets the data and file requests from the external computing devices. The file processor 252 then performs the file or data request based on the information it is given, and responds accordingly to the file or data request. The network file system is maintained and operated on solely by the file processor 252 of the SAN system 250. All accesses, writes, reads, and requests for information on any files and/or data under the network file system is handled by the SAN system 250, and in particular the file processor 252.

The file processor 252 keeps track of all the stored files and/or data stored on the media devices 257, 258, and 259. The file processor 252 maintains and presents a file system view of the stored data and/or files to the computing devices 210, 220, and 230 over the remainder of the SAN system 250 and the SAN network media 264. The computing devices 210, 220, and 230, when accessing or inquiring about portions of the network file system, perform these functions by requesting them through the file processor 252 of the SAN system 250.

The SAN system 250 allows access to the files and/or data stored in its storage media, and actually performs all the function of a file system to the attached computing devices 210, 220, and 230. Opening, closing, reading, and writing of data to files and of files themselves actually look and perform like a normal file system to the attached computing devices 210, 220, and 230. These actions are transparent to the computing devices. As such, the SAN system 250 acts and performs as a file system to the rest of the computing devices connected to the file processor 252. Also, from the perspective of the computing devices, each computing device can access and view the data and/or files stored by the file processor 252 of the SAN system 250 as part of a large, monolithic file system.

A client component 262 and a media component 260 can be part of the SAN system 250. These components work in conjunction with other components present in the network environment, including the file processor 252 itself, to make up a network backup and retrieval system for the computer network 290.

In an embodiment of the present invention, the file processor 252 works in conjunction with the management component 202, the media component 260, and the client component 262 for archival type requests, such as those concerned with backup, retrieval, and restoration purposes. The media component 260 acts in conjunction with the management component 202 and/or the client component 262 in a backup and retrieval operation with regards to the network files as stored on the SAN 250.

The management component 202 could, for example, initiate a full backup of the network file system as stored and managed on the SAN system 250. This could be initiated through the network link 270 directly to the client component 262, bypassing the SAN link 264.

Or, the management component 202 could initiate the action through any of the computing devices 210, 220, or 230. This initiation may take place either in a direct request to the SAN system 250 or indirectly to the components 260 and 262 through such methods such as data encapsulation and data bridging. Or, the initiation could be a special file memory request to the SAN system 250, which the file processor 252 interprets to be a particular backup and retrieval instruction.

It may also be possible that the client component 262 requests the backup itself, independently of the media component 260. In either event the client component 262 would manage the functions associated during the backup with the host system, in this case the SAN

system 250, such as determining the actual files or data to backup, the level of backup, and other such client machine specific determinations. The data and/or files that need to be backed up would be made available from the network file storage media 257, 258, and 259, wherein the client component 262 turns control over to the media component 260. The media component 260 would then direct the physical storage of the data and/or files on the network file system from the storage media 257, 258, or 259, as the case may be, and onto the library storage media 275. The media component 260 could then perform the indexing functions on the archived data and/or files.

It should be noted that the backup could take several forms. A backup could target data and files on a sector or block write basis, or could be used in a file basis.

In the case of an incremental backup, for example, only those blocks or files that have been altered would be stored for backup and retrieval purposes. In the case of a differential backup, only those changed blocks as contained within an altered file would be stored. Or, other criteria, such as file size, can be used to determine a hybrid backup strategy wherein both files and blocks are saved, depending on the criteria employed and the state of the data and/or files as they exist on the SAN system.

In a restore-type archival operation, a similar method would be employed. Either the media component 260 or the client component 262 may request a restore. In either case, the client component 262 would then perform the managerial tasks associated with the request, as described earlier. Control would then pass to the media component 260 to physically perform the extraction of the stored or archived data and/or files from the library media 275. The client component 262 would then forward the retrieved data and/or files to the requesting device.

Should the requesting device be the SAN system itself, the client component 262 would forward the retrieved data to the SAN system 250, wherein the SAN system 250 could write the data out to the appropriate storage media 257, 258, or 259.

Or, the client component 262 could forward the retrieved data and/or files to the management component 202, wherein the management component 202 routes the requested data and/or files to the particular computing device.

Alternately, the computing device 220 running the client component 222 may request a restore or other archival request for its attached memory media device 214 through the client component 222. The media component 260 could be contacted either as a special media access request to the SAN system 250, or it could access the media component 260 through such methods as data encapsulation over the SAN network 264. Once contacted, if the request was for retrieval or a restore, the media component 260 would collect the appropriate data and/or files and relay the retrieved data and/or files to the computing device 220 through a communication with the SAN system 250. This return communication could be in the form of a SAN communication of a network type file or data, or it could employ the use of data encapsulation or data bridging for the transmittal of the retrieved information.

If the request from the client component 222 is for archiving a file, block, or set of either of the two, the media component 260 could acknowledge the request either directly through a SAN type message from the SAN system 250, or by encapsulating the response in a SAN type message. The client component 222 running on the computing device 220 would then direct the appropriate data or files from the memory media 214 to the media component 260. This again may take place either through a special access protocol recognizable by the SAN system 250 and

redirected to the media component 260, or through encapsulating the data sent over the SAN link 264 from the computing device as a SAN-formatted message directed to the media component 260. It should be noted that the management component 202 running on a different computing device could also initiate a backup and retrieval request by the client component 222 through the network 270.

Turning now the computing device 230, the computing device 230 is running a client component 232 that manages its archiving needs. The computing device 230 is not in direct contact with the media component 260 operating on the library storage media 275. A request for an archival action such as retrieval, a restoration, or a backup is made by the client component 232. This request can be initiated either by the management component 202 or by the client component 232 itself. The client component 232 then coordinates and determines the scope of the backup and retrieval request, and accordingly acts to notify the media component 260. This may be accomplished either by a direct request to the SAN system 250 over the SAN link 264 acting as a request for a local backup and retrieval request, in which the SAN system 250 coordinates the backup and retrieval request. Or, this may be accomplished by the routing of a message directly for the media component 260 through use of data encapsulation via the SAN system 250.

If the request is a request for a backup, the client component 232 could then communicate the files and/or data to be archived to the media component 260 in a similar manner. The media component 260 would then perform the requested backup to the media library 275.

If the request is one for a retrieval or restoration, the media component 260 would extract the requested data from the media library 275 and route the data back to the client component

232 which would be responsible for the placement of the data on the computing system 230. This outbound messaging may be accomplished either by direct communication through the SAN system 250, or may be by indirect methods, such as data encapsulation from the media component 260 or the use of data bridging techniques.

Fig. 3 is schematic block diagram of the interaction of the library media of Fig. 2 with the SAN system. As shown, a library media 310 controlled by a media component 320 may comprise a number of different storage media, or may just comprise one. In Fig. 3, the library media 310 comprises a fast, alterable random access device 312, a fast, non-alterable random access device 314, a serial device 316, a slow, alterable random access device 318, and a slow, non-alterable random access device 319.

An example of the fast, alterable random access device 312 includes various magnetic media, such as a disc drive, that could include multiple writing surfaces. An example of the fast, non-alterable random access device 314 includes a multi disc optical system. An example of the slow, alterable random access device 318 includes jukeboxes containing disc drive cartridges. An example of the slow, non-alterable random access device 314 includes jukeboxes containing optical discs. An example of the serial device 316 could include a magnetic tape cartridge jukebox.

The media component 320 would control the placement of files, sectors, and other backup and retrieval information on the appropriate library media. This placement could be controlled according to the parameters of the backup, such as proximity in date, or whether the archived data is alterable in the archived form. Other parameters to consider could be the

relative frequency of requests to the data or to importance of the data as determined by a client component or a management component directing those parameters.

Thus, in the case of differential backups, portions of the archived file may reside across several different media. Older portions may be contained in the device 314, while newer updated versions of that block may be contained in the device 312. Portions that have not changed may still be in other library devices.

In view of the above detailed description of the present invention and associated drawings, other modifications and variations will now become apparent to those skilled in the art. It should also be apparent that such other modifications and variations may be effected without departing from the spirit and scope of the present invention as set forth in this specification.

044463.0024 AUSTIN 188229 v2